

The rise of enterprise search has put an increasing focus on searching ever broader collections of content and documents within organisations.

While enterprise search projects generally start with simple intranet search, attention quickly moves to searching document management systems, collaboration tools, business systems and fileshares.

Underpinning this work is the belief (or hope) that business value will be delivered to users by deploying a more extensive search tool.

Unfortunately it is often the case that searching more is not better than searching less.

This briefing will look at some of the challenges involved in implementing enterprise search, and provide practical tips on how to proceed.

Relevance and value

The fundamental goal of any search tool is to provide users with useful and relevant search results.

Within the enterprise, this means finding valuable information across the many different repositories, sources and systems.

The difficulty is that increasing the amount of information being searched almost always reduces the relevance of search results. Once called the ‘Altavista effect’, this was seen in the millions of hits generated for any set of terms entered into that search engine.

This is equally significant within an enterprise, and the challenge is to maintain (or improve) relevance as the volume of information grows.

Consideration also needs to be given to user needs and expectations. What types of queries are being entered into the enterprise search, and what sorts of results are expected?

Gaining a deeper understanding of these questions helps to shed light on what to search, and how.

A content problem

In many cases, enterprise search is deployed to help address the problems of fragmented information repositories, uneven content quality, and difficulties in finding required information.

These are, however, content problems rather than search problems. Deploying a powerful search often does nothing more than shine a very bright light on the underlying content issues in the enterprise.

While good search can help to mitigate these problems, it cannot resolve them. It is also time consuming and costly to set up search tools to handle poor quality content.

Enterprise search is therefore best targeted at well-managed information repositories within the organisation. Only then should incremental steps be taken to push out into less managed environments.

For example, enterprise search is almost invariably wasted when searching fileshares, as these chaotic collections of documents are the very problem that led to replacement platforms being put in place (such as intranets or document management systems).

Determining relevance

As the scope of enterprise search grows, it is also difficult to determine relevance. When a single search covers intranet pages, files in a document management system, and records out of the staff directory, which gets listed first?

This is a fundamental challenge for enterprise search, and one that needs to be well understood if useful results are to be returned to users.

Taking a pragmatic approach

While enterprise search is difficult, it is also valuable, and a pragmatic approach should be taken:

- Implement search usage reports to identify common and failed searches.
- Use this (and other) information to build a stronger understanding of what staff are looking for, and where this is likely to be found.
- Incrementally add repositories, testing on real data to find out whether they will increase or decrease relevance.
- Hide the details from users, maintaining a simple and usable interface that ‘just works’ regardless of what is being searched.